

The Fixed Initial Credit Problem for Partial-Observation Energy Games is ACK-complete

Guillermo A. Pérez*

Département d'Informatique, Université Libre de Bruxelles
gperezme@ulb.ac.be

November 20, 2016

Abstract

In this paper we study two-player games with asymmetric partial observation and an energy objective. Such games are played on a weighted automaton by Eve, choosing actions, and Adam, choosing a transition labelled with the given action. Eve attempts to maintain the sum of the weights (of the transitions taken) non-negative while Adam tries to do the opposite. Eve does not know the exact state of the game, she is only given an equivalence class of states which contains it. In contrast, Adam has full observation. We show the fixed initial credit problem for these games is ACK-complete.

1 Introduction

Energy games are two-player quantitative games of infinite horizon played on finite weighted automata. The game is played in rounds in which one player, Eve, chooses letters from the automaton's alphabet whilst Adam, the second player, resolves non-determinism. The initial configuration of the game is determined by an automaton, an initial state, and an *initial credit* for Eve. The goal of Eve in an energy game is to keep a certain resource from being depleted. More specifically, she wins if, for every round, the sum of the weights of the transitions traversed so far plus her initial credit is non-negative. Adam has the opposite objective: witnessing a negative value.

Quantitative games, in general, are useful models for the interaction of the controller and its environment in open reactive systems. Furthermore, synthesizing controllers in this setting reduces to computing a winning strategy for one of the players in the corresponding game. Energy games in particular are useful for systems in which one is interested in the use of bounded resources such as power or fuel [CdAHS03, BFL⁺08].

Two decision problems for energy games have been studied by the formal verification community: the *fixed initial credit* and *unknown initial credit problems*. The former asks whether, given a fixed initial credit for Eve, she has a strategy which ensures all plays consistent with it are winning. The latter is more ambitious in that it asks whether there exists some initial credit for which the same question has a positive answer. It is known that if Eve has a winning strategy in an energy game, she also has a memoryless winning strategy. Furthermore, in order to win, Eve essentially has to ensure staying in cycles with non-negative (total) weight. Using these two facts, one can show the unknown initial credit reduces in polynomial time to the fixed initial problem. (If there is some initial credit for which Eve wins, then nw_{\max} should suffice—where n is the number of states and w_{\max} denotes the maximum absolute value of a transition weight in the automaton.) It is also known that the fixed initial credit problem is log-space equivalent to the threshold problem for *mean-payoff games* [BFL⁺08]. In a mean-payoff game, the objective of Eve consists in maximizing the limit (inferior) of the averages of the running sum of transition weights observed along an infinite play. Determining if Eve can ensure at least a given mean-payoff value is not known to be decidable in polynomial time. However, as they are known to be positionally determined [ZP96], the latter problem can be decided in NP and in coNP. This places

*Author supported by an F.R.S.-FNRS fellowship.

energy games in a very special class of problems which are known to be solvable in $\text{NP} \cap \text{coNP}$ and for which no polynomial-time algorithm has been discovered.¹

Multi-dimensional energy games have been studied in, amongst other works, [CDHR10] and more recently in [JLS15]. These are the natural generalization of energy games played on singly-weighted automata to automata with vector weights on their transitions. They are relevant to the synthesis of reactive controllers sensitive to the usage of multiple resources. In this setting, it has been shown that the fixed initial credit problem is 2EXPTIME-complete while the unknown initial credit problem is simpler, namely, coNP-complete.

In the present work we focus on another generalization, namely, *energy games with partial observation*. Such games are, once more, played on finite singly-weighted automata. The difference between classical energy games and partial-observation energy games is that, in the latter, a partition of the states of the automaton into *observations* is also given as part of the input. The game is then modified so that, after every round, Eve is only informed of the observation of the successor state chosen by Adam (hence the partial observability). Energy games with partial observation were initially studied in [DDG⁺10] and [HPR14]. They capture the fact that controllers in open reactive systems have limited capabilities, e.g. a finite number of sensors with limited precision. Two results from [DDG⁺10] are of particular interest to us. First, the unknown initial credit problem for partial-observation energy games was shown to be undecidable by reduction from the halting problem for Minsky machines. Second, decidability of the fixed initial credit problem was established by describing a reduction to finite safety games.

Contributions In this work we refine the upper and lower bounds for the fixed initial credit problem for partial-observation energy games. Specifically, we show the size of the safety game used in the algorithm from [DDG⁺10] is at most Ackermannian with respect to the size of the input game. We then describe how the Minsky machine simulation used to show undecidability of the unknown initial credit problem in [DDG⁺10] can be modified to show ACK-hardness of the fixed initial credit problem. This establishes ACK-completeness of the problem.

2 Preliminaries

Games A *weighted game with partial observation* (or just game, for short) is a tuple $(Q, q_0, \Sigma, \Delta, w, \text{Obs})$ where Q is a non-empty finite set of states, $q_0 \in Q$ is the initial state, Σ is a finite alphabet of actions or symbols, $\Delta \subseteq Q \times \Sigma \times Q$ is a total transition relation, $w : \Delta \rightarrow \mathbb{Z}$ is a weight function, and $\text{Obs} \subseteq \mathcal{P}(Q)$ is a partition of Q into *observations*. If $\text{Obs} = \{Q\}$ we say the game is *blind*, if $\text{Obs} = \{\{q\} \mid q \in Q\}$ we say it is a *full-observation* game. For $s \subseteq Q$ and $\sigma \in \Sigma$, denote by $\text{post}_\sigma(s) := \{q' \in Q \mid \exists q \in s \wedge (q, \sigma, q') \in \Delta\}$ the set of σ -successors of s . Also, let w_{\max} denote the maximum absolute value of a transition weight in the automaton.

Unless otherwise stated, in what follows we consider a fixed $G = (Q, q_0, \Sigma, \Delta, w, \text{Obs})$.

Plays A *play* in G is an infinite sequence $\pi = q_0\sigma_0q_1\sigma_1\ldots$ such that $(q_i, \sigma_i, q_{i+1}) \in \Delta$, for all $i \geq 0$. For a play π and integers $0 \leq i \leq j$, we denote by $\pi[i..j]$ the infix $q_i\sigma_i\ldots a_{j-1}q_j$ of π . The set of plays in G is denoted by $\text{Plays}(G)$ and the set of prefixes of plays ending in a state is written $\text{Prefs}(G)$. The unique observation containing state q is denoted by $\text{obs}(q)$. We extend $\text{obs}(\cdot)$ to plays and prefixes in the natural way. For instance, we obtain the *observation sequence* $\text{obs}(\pi)$ of a play π as follows: $\text{obs}(q_0)\sigma_0\text{obs}(q_1)\sigma_1\ldots$

Objectives An objective in a game corresponds to a set of “good” plays.

The *energy level* of a play prefix $\pi = q_0\sigma_0\ldots q_n$ is $\text{EL}(\pi) = \sum_{i=0}^{n-1} w(q_i, \sigma_i, q_{i+1})$. The *energy objective* is parameterized by an initial credit $c_0 \in \mathbb{N}$ and is defined as:

$$\text{PosEn}(c_0) := \{\pi \in \text{Plays}(G) \mid \forall i > 0 : c_0 + \text{EL}(\pi[0..i]) \geq 0\}.$$

In other words, the energy objective asks for the energy level of a play never to drop below 0 when starting with energy level c_0 .

¹Other games known to be in the same class are discounted-sum and simple stochastic games [ZP96]. In fact, all games mentioned in this work have further been shown to be in $\text{UP} \cap \text{coUP}$ [Jur98] and are therefore unlikely to be NP-complete.

We will also make use of the *safety* objective, defined relative to a set of *safe states* $W \subseteq Q$:

$$\text{Safe}(W) := \{q_0\sigma_0q_1 \cdots \in \text{Plays}(G) \mid \forall i \geq 0 : q_i \in W\}.$$

Intuitively, the objective asks that Eve make sure the play never visits states outside of W . Note that in safety games the weight function is not needed.

Strategies A *strategy for Eve* is a function $\lambda : \text{Pref}(G) \rightarrow \Sigma$. A strategy λ for Eve is *observation-based* if for all prefixes $\pi, \pi' \in \text{Pref}(G)$, if $\text{obs}(\pi) = \text{obs}(\pi')$ then $\lambda(\pi) = \lambda(\pi')$. A prefix (or play) $\pi = q_0\sigma_0 \dots$ is consistent with a strategy λ for Eve if $\lambda(\pi[0..i]) = \sigma_i$, for all $i \geq 0$. We say a strategy λ for Eve is a *winning strategy for her* in a game with objective Ω if all plays consistent with λ are in Ω .

We do not formalize the notion of strategy for Adam here. Intuitively, given a play prefix and an action $\sigma \in \Sigma$, he selects a σ -successor q of the current state and reveals $\text{obs}(q)$ to Eve.

Problem 1 (Fixed initial credit problem). *Given a game G and an initial credit c_0 , decide whether there exists a winning observation-based strategy for Eve for the objective $\text{PosEn}(c_0)$.*

3 Upper bound

The fixed initial credit problem was shown to be decidable in [DDG⁺10]. To do so, the problem is reduced to determining the winner of a safety game played on a finite tree whose nodes are functions which encode the *belief* of Eve in the original game. The notion of belief corresponds to the information Eve has about the current state (at any round) of a partial-observation game. In this particular case, the belief of Eve is defined by a prefix $\pi = q_0\sigma_0 \dots \sigma_{n-1}q_n$ with $o = \text{obs}(q_n)$. It corresponds to the subset of states $s \subseteq o$ which are reachable from q_0 via a prefix π' with the same observation sequence as π , i.e. $\text{obs}(\pi) = \text{obs}(\pi')$, together with the energy levels of all prefixes ending in q for all $q \in s$. Note that Eve only really cares about the minimal energy levels of prefixes ending in states from s . This information can thus be encoded into functions.

In this section we will first formalize the construction described above. We will then give an alternative argument (to the one presented in [DDG⁺10]) which proves that the constructed tree is finite. The latter goes via a translation from functions to vectors and prepares the reader for the next result. Finally, we will define an Ackermannian function and show that the size of the tree is at most the value of the function on the size of the input game.

Throughout the section we consider a fixed partial-observation energy game $G = (Q, q_0, \Sigma, \Delta, w, \text{Obs})$ and fixed initial credit $c_0 \in \mathbb{N}$.

3.1 Reduction to safety game

Belief functions We define the set of *belief functions* of Eve as $\mathcal{F} := \{f : Q \rightarrow \mathbb{Z} \cup \{\perp\}\}$. The *support* of a function $f \in \mathcal{F}$ is the set $\{q \in Q \mid f(q) \neq \perp\}$. A function $f \in \mathcal{F}$ is said to be *negative* if $f(q) < 0$ for some $q \in \text{supp}(f)$. The initial belief function f_0 has support $\{q_0\}$ and $f_0(q_0) = c_0$. Given two functions $f, g \in \mathcal{F}$ we define the order $f \preceq g$ to hold if $\text{supp}(f) = \text{supp}(g)$ and $f(q) \leq g(q)$ for all $q \in \text{supp}(f)$. Additionally, for $\sigma \in \Sigma$ we say g is a σ -*successor* of f if $\exists o \in \text{Obs} : \text{supp}(g) = \text{post}_\sigma(\text{supp}(f)) \cap o$ and $g(q) = \min\{f(p) + w(p, \sigma, q) \mid p \in \text{supp}(f) \wedge (p, \sigma, q) \in \Delta\}$ for all $q \in \text{supp}(g)$. Intuitively, if Eve has belief function f and she plays σ , then if Adam reveals observation o to her as the observation of the new state of the game, she now has belief function g .

Function-action sequences For a function-action sequence $s = f_0\sigma_0f_1 \dots \sigma_{n-1}f_n$ we will write f_s to denote f_n , i.e. the last function of the sequence. Let S be the smallest subset of $(\mathcal{F} \cdot \Sigma)^*\mathcal{F}$ containing f_0 and $s \cdot \sigma \cdot f$ if $s \in S$, f is a σ -successor of f_s , and it holds that: (a) f_s is not negative and (b) $f_{s'} \not\preceq f_s$ for all proper prefixes s' of s . The desired full-observation safety game is then $H = (S, f_0, \Sigma, E, W)$ where

- the transition relation $E \subseteq S \times \Sigma \times S$ contains triples (s, σ, s') where $s' = s \cdot \sigma \cdot f_{s'}$, and
- the safe states are $W = \{s \in S \mid f_s \text{ is not negative}\}$.

In order for E to be total, we add self-loops (s, σ, s) for any $s \in S$ without outgoing transitions.

Lemma 1 (From [DDG⁺10]). *There is a winning observation-based strategy for Eve for the objective $\text{PosEn}(c_0)$ in G if and only if there is a winning strategy for Eve in the safety game H .*

3.2 Showing the safety game is finite

Henceforth, let $H = (S, f_0, \Sigma, E, W)$ be the safety game constructed from the partial-observation game G and initial credit c_0 we have fixed for all of Section 3.

In the sequel, it will be useful to consider vectors instead of functions. We will therefore define an encoding of belief functions into vectors. Formally, let us fix two bijective mappings $\alpha : \{1, \dots, |Q|\} \rightarrow Q$ and $\beta : \mathcal{P}(Q) \rightarrow \{1, \dots, 2^{|Q|}\}$. (The latter two mappings essentially corresponding to fixing an ordering on Q and the set of subsets of Q .) For a belief function f , we will define a vector $\vec{f} \in \mathbb{Z}^{|Q|+2}$ which holds in its i -th dimension the value assigned by f to state $\alpha(i)$ in its support. For technical reasons (see Lemma 5), if state $\alpha(i)$ is not part of the support of f , we will use as place-holder the minimal value assigned by f to any state. Additionally, we use two dimensions to identify uniquely the support set of f . More formally, the vector \vec{f} is

$$(2^{|Q|} - \beta(\text{supp}(f)), \beta(\text{supp}(f)), \gamma \circ \alpha(|Q|), \dots, \gamma \circ \alpha(1))$$

where $\gamma(q)$ is $f(q)$ if $q \in \text{supp}(f)$ and $\min\{f(q) \mid q \in \text{supp}(f)\}$ otherwise.

It follows directly from the above definitions that two belief functions being \preceq -comparable is sufficient and necessary for their corresponding vectors to be \leq -comparable.² More formally,

Lemma 2. *For belief functions $f, g \in \mathcal{F}$ we have that $f \preceq g$ if and only if $\vec{f} \leq \vec{g}$.*

Using the above Lemma we can already argue that the safety game H is finite. Suppose, towards a contradiction, that H is infinite. By König's Lemma, there is an infinite function-action sequence $s = f_0 \sigma_0 f_1 \dots$ such that for all $i \geq 0$ we have: (i) f_{i+1} is a σ_i -successor of f_i , (ii) f_i is not negative, and (iii) $f_j \not\preceq f_i$, for all $0 \leq j < i$. Now, let us consider the vector sequence $v = \vec{f}_0 \vec{f}_1 \dots$. Note that the two dimensions used to represent the support of the function cannot be negative. Hence, together with condition (ii) above, it follows that $\vec{f}_i \in \mathbb{N}^{|Q|+2}$ for all $i \geq 0$. That is, the vectors have no negative integers. Further, from (iii) together with Lemma 2 it follows that there are no distinct vectors \vec{f}_i, \vec{f}_j in the sequence v such that $\vec{f}_i \leq \vec{f}_j$. Thus, we get a contradiction with Dickson's Lemma—which states that every infinite sequence of vectors of natural numbers has two distinct \leq -comparable elements.

Proposition 1 (From [DDG⁺10]). *The game H has finite state space.*

3.3 An Ackermann bound on the size of the safety game

The argument we have presented, to show the safety game is finite, carries the intuition that any function-action sequence from H should eventually end in a negative function or a function which is bigger than another function in the sequence (w.r.t. \preceq). How long can such sequences be? Using the relation between functions and vectors that we have established (and formalized in Lemma 2) we can apply results of Schmitz et al. [FFSS11, SS12, Sch16] which have been formulated for sequences of vectors of natural numbers. Intuitively, the bound they provide is based on “how big the jump is” from each vector in the sequence to the next. This last notion is formalized in the following definition.

Definition 1 (Controlled vector sequence). *A vector sequence $\mathbf{a}_0 \mathbf{a}_1 \dots \in (\mathbb{N}^d)^*$ is t -controlled by a unary increasing function $\kappa : \mathbb{N} \rightarrow \mathbb{N}$ if $|\mathbf{a}_i|_\infty < \kappa(t + i)$ for all $i \geq 0$.³*

We will now define a hierarchy of sets of functions. Our intention is to determine at which level of this hierarchy we can find a function which can be said to control vector sequences induced by function-action sequences from H . This will allow us to find a function—also at a specific level of the hierarchy—that bounds the length of such sequences (see Lemma 6).

The fast-growing functions These can be seen as a sequence $(F_i)_{i \geq 0}$ of number-theoretic functions defined inductively below. [FW98]

$$F_0(x) := x + 1$$

$$F_{i+1}(x) := F_i^{x+1}(x) = \overbrace{F_i(F_i(\dots F_i(x) \dots))}^{x+1 \text{ times}}$$

²To be precise, \leq here denotes the product ordering on vectors of integers.

³For a vector $\mathbf{a} = (a_d, a_{d-1}, \dots, a_1)$, the infinity norm is the maximum value on any dimension, i.e. $\max\{a_i \mid 1 \leq i \leq d\}$.

The following Lemma summarizes some properties of the hierarchy.

Lemma 3 (From [FW98]). *For all $i \in \mathbb{N}$,*

- *for all $i \leq j$ and $0 \leq n \leq m$, $F_j(m) \geq F_i(n)$ and the latter is strict if the inequality between j and i or m and n is strict;*
- *F_i is primitive-recursive;*
- *F_i is dominated by F_{i+1} .⁴*

Furthermore, for all primitive-recursive functions f , there exists $i \in \mathbb{N}$ such that F_i dominates f .

We consider the following variant of the Ackermann function $F_\omega(x) := F_x(x)$. It is not hard to show that F_ω dominates all F_i —that is, for all $i \in \mathbb{N}$ —and, in turn, all primitive-recursive functions.

The Grzegorzcyk hierarchy We now introduce a sequence $(\mathfrak{F}_i)_{i \geq 2}$ of sets of functions. Using the i -th fast-growing function, we define the i -th level of the hierarchy [Wai70, Sch16] as follows:

$$\mathfrak{F}_i := \bigcup_{c \in \mathbb{N}} \text{FDTIME}(F_i^c(x)).$$

In other words, \mathfrak{F}_i consists of all functions $\mathbb{N} \rightarrow \mathbb{N}$ which can be computed by a deterministic Turing machine in time bounded by any finite composition of the function F_i . Note that, since F_2 is of exponential growth, we could restrict space instead of time or even allow non-determinism and obtain exactly the same classes.

The following property of the classes of functions from the hierarchy will be useful.

Lemma 4 (From [LW70, Sch16]). *For all $i \geq 2$, every $f \in \mathfrak{F}_i$ is dominated by F_j if $i < j$.*

We will now show how to control vector sequences induced by function-action sequences from H (following the Karp-Miller tree analysis from [FFSS11]).

Lemma 5. *For all non-negative function-action sequences $s \in S$, the corresponding vector sequence from $(\mathbb{N}^{|Q|+2})^*$ is $(c_0 + w_{\max} + |Q|)$ -controlled by $k(x) := 2^x + x^2$.*

Proof. Let us assume that $w_{\max} > 0$. (This is no loss of generality as the energy game is trivial otherwise.) For any sequence $s = f_0 \sigma_0 \dots \sigma_{n-1} f_n \in S$ we have that for all $0 \leq i \leq n$:

$$\begin{aligned} \|\vec{f}_i\|_\infty &\leq \max\{2^{|Q|}, c_0 + i \cdot w_{\max}\} \\ &\leq 2^{|Q|} + (c_0 + w_{\max} + i)^2 \\ &\leq 2^{|Q|+c_0+w_{\max}+i} + (|Q| + c_0 + w_{\max} + i)^2 \end{aligned}$$

which concludes the proof. \square

Note that the control function from Lemma 5 is at the second level of the Grzegorzcyk hierarchy. That is, $k \in \mathfrak{F}_2$, since F_2 is exponential. We can now apply the following tool.

Lemma 6 (From [FFSS11]). *For natural numbers $d, i \geq 1$, for all unary increasing functions $\kappa \in \mathfrak{F}_i$, there exists a function $L_{d,\kappa} : \mathbb{N} \rightarrow \mathbb{N} \in \mathfrak{F}_{i+d-1}$ such that $L_{d,\kappa}(t)$ is an upper bound for the length of non-increasing sequences from $(\mathbb{N}^d)^*$ that are t -controlled by κ .*

To conclude, we show how to bound the length of any sequence $s \in S$. By construction of H , s is a function-action non-increasing sequence $f_0 \dots f_n$ of non-negative functions—except for the last function, which might be negative. The vector sequence $\vec{f}_0 \dots \vec{f}_{n-1}$ is therefore non-increasing and, for all $0 \leq i < n$, the vector \vec{f}_i has dimension $|Q| + 2$ and contains only non-negative numbers. It follows from Lemmas 5 and 6 that the length of the vector sequence is less than $h(c_0 + |Q| + w_{\max})$, where $h = L_{|Q|+2,k} \in \mathfrak{F}_{|Q|+3}$. Hence, the length of s is bounded by $h(c_0 + |Q| + w_{\max}) + 1$. Let us write $|G| = |Q| + c_0 + w_{\max} + |\Delta| + |\text{Obs}|$. We thus have that $h(|G|) + 1$ bounds the length of all $s \in S$. Clearly then, the size of S is at most $|\Delta|^{h(|G|)+1}$. More coarsely, we have that $2^{(h(|G|)+1)^2}$ bounds the size of S . It follows from Lemmas 3–6 that the latter bound is primitive recursive for all fixed G . Since F_ω dominates all primitive-recursive functions, we conclude $|S|$ is $\mathcal{O}(F_\omega(|G|))$. As safety games are known to be solvable in linear time with respect to the size of the game graph (see, e.g., [AG11]), the desired result then follows from Lemma 1.

Theorem 1. *The fixed initial credit problem is decidable in Ackermannian time.*

⁴For two functions $f, g : \mathbb{N} \rightarrow \mathbb{N}$, we say g dominates f if $g(x) \geq f(x)$ for all but finitely many $x \in \mathbb{N}$.

4 Lower bound

In the sequel we will establish Ackermannian hardness of the fixed initial credit problem, thus giving a negative answer to the question of whether the problem has a primitive-recursive algorithm. This question is of particular interest in light of recent work by Jurdziński et al. [JLS15] in which it is shown the same problem is 2EXPTIME-complete for multi-dimensional games with full observation.

To begin, we will formally define the ACK complexity class—using the hierarchies of functions introduced in the previous section. We will then adapt the translation from Minsky machines to partial-observation energy games presented in [DDG⁺10] (to argue the unknown initial credit problem is undecidable) and reduce the existence of a halting run with bounded counter values in the original machine to the fixed initial credit problem in the constructed game. Finally, we will describe how to make sure the bound on the counters is Ackermannian (without explicitly computing the Ackermann function during the reduction).

The complexity class We adopt here the definition proposed by Schmitz [Sch16] for the class of Ackermannian decision problems:

$$\text{ACK} := \bigcup_{g \in \mathfrak{F}_{<\omega}} \text{DTIME}(F_\omega(g(n)))$$

where $\mathfrak{F}_{<\omega} := \bigcup_{i \in \mathbb{N}} \mathfrak{F}_i$. Note that we allow ourselves any kind of primitive-recursive reduction. It was shown in [Sch16] that: for any two functions $g, f \in \mathfrak{F}_{<\omega}$, there exists p in $\mathfrak{F}_{<\omega}$ such that $f \circ F_\omega \circ g$ is dominated by $F_\omega \circ p$. It follows the distinction between time-bounded and space-bounded computations is actually irrelevant here since F_2 is already of exponential growth.

4.1 Minsky machine simulation

A 2-counter *Minsky machine* (2CM) consists of a finite set of control states Q , initial and final states $q_I, q_F \in Q$, a set of two counters C , and a finite set of instructions which act on the counters. Namely, inc_k increases the value of counter k by 1, dec_k decreases the same value by 1. Additionally, $0?_k$ serves as a *zero-check* on counter k which blocks if the value of counter k is not equal to 0. More formally, the transition relation δ contains tuples (q, ι, k, q') where $q, q' \in Q$ are source and target states respectively, ι is an instruction from $\{inc, dec, 0?\}$ which is applied to counter $k \in C$. We focus here on deterministic 2CMs, i.e. for every state $q \in Q$ either

- δ has exactly one outgoing transition, which is an increase instruction, i.e. (q, ι, \cdot, \cdot) with $\iota = inc$;
- or
- δ has two transitions: a decrease (q, dec, k, \cdot) and a zero-check $(q, 0?, k, \cdot)$ instruction.

We write $|M|$ instead of $|Q|$ to denote the *size of* M . A *configuration* of M is a pair (q, v) of a state $q \in Q$ and a valuation $v : C \rightarrow \mathbb{N}$.⁵ A *run* of M is a finite sequence $\rho = (q_0, v_0)\delta_0 \dots \delta_{n-1}(q_n, v_n)$ such that $q_0 = q_I$, $v_0(k) = 0$ for all $k \in C$, and v_{i+1} is the correct valuation of the counters after applying δ_i to v_i for all $1 \leq i \leq n$. A run $(q_0, v_0)\delta_0 \dots \delta_{n-1}(q_n, v_n)$ is *halting* if $q_n = q_F$ and it is *m-bounded* if $v_i(k) \leq m$ for all $0 \leq i \leq n$ and all $k \in C$.

Problem 2 (*f*-Bounded halting problem). *Given a 2CM M , decide whether M has an $f(|M|)$ -bounded halting run.*

This bounded version of the halting problem is decidable. However, if $f = F_\omega$, then it is ACK-complete [SS12, Sch16].

Lemma 7. *A 2CM M has an $f(|M|)$ -bounded halting run if and only if it has an $f(|M|)$ -bounded halting run of length at most $|M|(f(|M|))^2$.*

Proof. We focus on the only non-trivial direction. If no $f(|M|)$ -bounded run of the machine does reach q_F in at most $|M|(f(|M|))^2$ steps, then we have two possibilities. It could be the case that the machine has a counter whose value goes above $f(|M|)$ before reaching q_F . Clearly, M has no $f(|M|)$ -bounded halting run in this case. Otherwise, the machine must have repeated at least one configuration. Since the machine is deterministic, this means it will never halt (i.e. it will cycle without reaching q_F). \square

For the rest of this section, let us consider a fixed 2CM $M = (Q, q_I, q_F, C, \delta)$.

⁵Note that we consider the variant of Minsky machines which guards all decreases with zero-checks. Hence, all counters will have only non-negative values at all times.




Figure 1: Gadget which ensures the first letter played by Eve is #.




Figure 2: Gadget which ensures σ_1 is followed by σ_2 .

Energy game for the id-bounded halting problem We will now construct a blind energy game G_M such that M has an $|M|$ -bounded halting run if and only if there is a winning observation-based strategy for Eve in G_M given initial credit $c_0 = 0$, i.e. we reduce from the id-bounded halting problem where $\text{id}(x) = x$ is the identity function. Her winning observation-based strategy will be to perpetually simulate the halting run of the machine. We set the alphabet Σ of G_M to be the set of transitions of M plus a fresh symbol #, that is $\Sigma = \delta \cup \{\#\}$. The game G_M starts with a non-deterministic transition into one of several gadgets we will describe now. The weight of the transition into each gadget is shown (labelling the initial arrow) in the Figures for the gadgets. Each gadget will check the sequence of letters played by Eve has some specific property, lest some play will get a negative energy level. Since the game is blind, Eve will not know which gadget has been chosen and will therefore have to make sure her strategy (infinite word) has all the properties checked by each gadget.

The gadget depicted in Figure 1 makes sure that Eve plays # as her first letter. Indeed, if she plays a strategy which does not comply then there is a play which will end up in the -1 loop and thus gets a negative energy level. If she plays #, all plays in this gadget go to the 0 loop and will never have a negative energy level. To make sure that after # Eve plays the first transition from M , and then the second, \dots , we have $|\delta| + 1$ instances of the gadget from Figure 2. (Additionally, after having played a transition leading to q_F in M —i.e. of the form $(\cdot, \cdot, \cdot, q_F)$ —she must play # once more.) The intuition behind the gadget is simple: if she violates the order of the sequence of transitions then there will be a play consistent with her strategy which, in the corresponding gadget, reaches the -1 loop. If she plays the letters in the correct sequence then all plays in all gadgets can only stay in the initial state or go to the upper 0 loop.

To verify that Eve plays the letter # infinitely often, symbolizing the start of a new simulation of M every time, the game G_M includes the gadget shown in Figure 3. If Eve plays the letter # infinitely often and within $|M|^3$ rounds of each other, all the plays in the gadget will take the only negatively-weighted transition in the gadget at most $|M|^3$ times. Hence, all plays in this gadget will never have a negative energy level. If, however, Eve plays in any other way (eventually stopping with the letter # or taking too long to produce it) then there will be a play which reaches the middle state of the gadget and take the negative transition enough times to get a negative energy level. In Figure 4 we can see a very simple gadget which is instantiated for each $k \in C$. Its task is to ensure Eve does not play a sequence of inc_k and dec_k which results in the value of k being larger than $|M|$.

Finally, we make sure that Eve does not play a transition which executes a zero-check incorrectly. To do so, for each $k \in C$ we add to G_M an instance of the gadget from Figure 5. The intuition of how the gadget works is as follows. If Eve executes a $0?_k$ instruction when the counter value is positive (a *zero-cheat*), then there will be a play which goes to the top-right corner of the gadget and simulates the inverse of the intended operations on k and moves back to the initial state on the zero-cheat. This play thus far has an energy level of at most $|M|^3 - 1$. If Eve executes a dec_k instruction when the counter value is 0 (a *positive cheat*), then there will be a play which goes to the bottom-left corner of the gadget and simulates the operations on k and moves back to the initial state on the positive cheat. The latter




Figure 3: Gadget which ensures that Eve restarts her simulation of M infinitely often and that each simulation is of length at most $|M|^3$.




Figure 4: Gadget which ensures that Eve respects the bound on the counters.

play also has an energy level of at most $|M|^3 - 1$. It follows that if Eve cheats more than $|M|^3$ times, there will be a play in the gadget with negative energy level. If, however, she correctly simulates the zero checks, then a play can forever stay in the initial state—in which case it will never have a negative energy level—or it can move to the top-right or bottom-left corner at some point. There, if Eve is simulating a finite halting run ρ , she will play $\#$ again in at most $|\rho| + 1$ steps. If the play is still in one those corners at that moment, then we know it moves to the bottom-right state and that it must have seen a weight of -1 at most $|\rho|$ times. Clearly, once there, the play can no longer have a negative energy level. If the play returned to the initial state before that then, since Eve is not cheating, the energy level of the play must have been at least $c_0 + |\mathcal{M}|^3$ and it must have seen a weight of -1 at most $|\rho|$ times.

Proposition 2. *M has an $|M|$ -bounded halting run if and only if there is a winning observation-based strategy for Eve in G_M given initial credit $c_0 = 0$.*

Proof. If M does have an $|M|$ -bounded halting run ρ , then we can assume that ρ has length at most $|M|^3$ (see Lemma 7). Eve can then play the blind strategy which corresponds to the infinite word $(\#\rho)^\omega$. Since this word satisfies all the constraints ensured by the gadgets in G_M , no play will ever have negative energy level.

Suppose M has no $|M|$ -bounded halting run. Eve cannot play a strategy which does not correspond to a valid run of M or there will be a play with negative energy level in gadgets 1, 2, or 5. Thus, let us assume she does simulate $|M|$ faithfully. We now consider two cases depending on whether $|M|$ has a halting run (which is, necessarily, not $|M|$ -bounded). If $|M|$ has a halting run which is not $|M|$ -bounded, then a play with negative energy level can be constructed in gadget 4. Similarly, if her simulation of $|M|$ stays $|M|$ -bounded but takes longer than $|M|^3$ steps (because it is not halting), the strategy will not be winning because of gadget 3. Hence, she has no observation-based winning strategy. \square

We will now generalize the reduction we just presented and use it to prove the announced ACK-hardness result. In short, we need to make sure the energy level of plays entering gadget 4 becomes $F_\omega(|M|)$ and the energy level of plays entering gadgets 3 and 5 becomes greater than $|M|(F_\omega(|M|))^2$. The way in which we propose to do so is to add an initial gadget to G_M which allows Eve to play a specific sequence of letters to get the required energy level—and no more—and then forces here into the simulation of M which we have already described. As a first step, we describe a way of computing F_ω using vectors.

4.2 Vectorial version of the fast-growing functions

We will now give a vectorial-based definition of F_k , for any $k \in \mathbb{N}$. Intuitively, we will use $k + 1$ dimensions to keep track of how many iterations of F_j (for $0 \leq j \leq k$) still need to be applied on the current intermediate value. More formally, for a vector $\mathbf{a} = (a_k, \dots, a_0) \in \mathbb{N}^{k+1}$ we set

$$\Phi(\mathbf{a}; x) = \Phi(a_k, \dots, a_0; x) := F_k^{a_k}(\dots F_1^{a_1}(F_0^{a_0}(x))).$$

It follows that $F_\omega(x) = \Phi(1, \overbrace{0, \dots, 0}^{x \text{ times}}; x)$. The following is the key property associated with Φ .

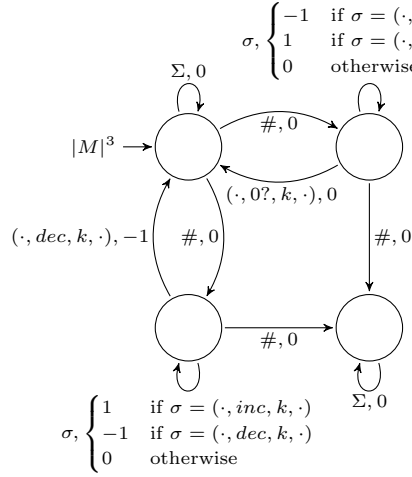


Figure 5: Gadget which ensures Eve correctly resolves the guarded decreases: executing a dec_k instruction only when $k > 0$ and executing the $0?_k$ instruction otherwise.

Lemma 8. For vectors $\mathbf{a}, \mathbf{b} \in \mathbb{N}^{k+1}$ and $x, y \in \mathbb{N}$, if $\mathbf{a} \leq \mathbf{b}$ and $x \leq y$ then $\Phi(\mathbf{a}; x) \leq \Phi(\mathbf{b}; y)$.

Proof. It is a direct consequence of the definition of Φ and Lemma 3. \square

We consider the following (family of) rewrite rules $N0$, $N1_j$, and $N2$:

$$\begin{aligned} \Phi(\mathbf{a}; x) &\rightarrow_{N0} x \\ \Phi(\dots, a_j + 1, a_{j-1}, \dots; x) &\rightarrow_{N1_j} \Phi(\dots, a_j, x + 1, \dots; x) \\ \Phi(a_k, \dots, a_0 + 1; x) &\rightarrow_{N2} \Phi(a_k, \dots, a_0; x + 1) \end{aligned}$$

For simplicity, denote the set of rewrite rules $\{N1_j \mid 0 < j \leq k\}$ by $N1$. Let us write $(\mathbf{b}, y) \rightsquigarrow_r (\mathbf{a}, x)$ and $(\mathbf{b}, y) \rightsquigarrow_N (\mathbf{a}, x)$ if rule r or, respectively, a sequence of the rules $N1$ and $N2$, can be applied to $\Phi(\mathbf{b}; y)$ to transform it into $\Phi(\mathbf{a}; x)$.⁶ We remark that $(\mathbf{b}, y) \rightsquigarrow_N (\mathbf{a}, x)$ implies \mathbf{a} is smaller than \mathbf{b} for the lexicographical order. It follows that the application of the rewrite rules always terminates.

Lemma 9. For any vector $\mathbf{b} \in \mathbb{N}^{k+1}$ and $x \in \mathbb{N}$, the set $\{\mathbf{a} \in \mathbb{N}^{k+1} \mid \exists y \in \mathbb{N} : (\mathbf{b}, y) \rightsquigarrow_N (\mathbf{a}, x)\}$ is finite.

Remark that rule $N1_j$ can be applied to $\Phi(\mathbf{a}; x)$ for any $0 < j \leq k$ as long as $a_j > 0$. We would like to argue that the “best” way to use $N1_j$, in order to obtain the highest possible final value, is to do so only if all dimensions $0 \leq \ell < j$ have value 0. Formally, let us write $(\mathbf{b}, y) \rightarrow_r (\mathbf{a}, x)$ if $(\mathbf{b}, y) \rightsquigarrow_r (\mathbf{a}, x)$ and, additionally, if $r = N1_j$ then it holds that $a_\ell = 0$ for all $0 \leq \ell < j$. We then say the application of the rewrite rule r was *proper*. Similarly, we write $(\mathbf{b}, y) \rightarrow_N (\mathbf{a}, x)$ if $\Phi(\mathbf{a}; x)$ can be obtained by proper application of rules $N1$ and $N2$ to $\Phi(\mathbf{b}; y)$.

Lemma 10. For all vectors $\mathbf{a}, \mathbf{b} \in \mathbb{N}^{k+1}$ and $x, y \in \mathbb{N}$, if $(\mathbf{b}, y) \rightarrow_N (\mathbf{a}, x)$ then $\Phi(\mathbf{b}; y) = \Phi(\mathbf{a}; x)$.

Proof. Follows directly from the definitions of Φ and the fast-growing functions, and proper application of rules $N1$ and $N2$. \square

The above result means that proper application of the rewrite rules to $\Phi(a_k, \dots, a_0; x)$ give us a correct computation of $F_k^{a_k}(\dots(F_0^{a_0}(x)))$. We will now argue that improper application of the rules will result in a smaller value.

Lemma 11. For all vectors $\mathbf{a}, \mathbf{b} \in \mathbb{N}^{k+1}$ and $x, y \in \mathbb{N}$, if $|\mathbf{b}|_\infty \leq x$ and $(\mathbf{b}, y) \rightsquigarrow_N (\mathbf{a}, x)$ then $\Phi(\mathbf{b}; y) \leq \Phi(\mathbf{a}; x)$.

⁶Since the rule $N0$ yields a single number, it cannot be the case that $(\mathbf{b}, y) \rightsquigarrow_{N0} (\mathbf{a}, x)$.




Figure 6: Pumping gadget with only transitions for $\sigma = N0$ shown for states $\{\chi\} \cup \{\alpha_i \mid 0 \leq i \leq m\}$.

Proof. If $(\mathbf{b}, y) \rightarrow_N (\mathbf{a}, x)$ then the result follows by applying Lemma 10. If this is not the case, the sequence of rules applied to $\Phi(\mathbf{b}; y)$ to obtain $\Phi(\mathbf{a}; x)$ includes at least one improperly applied rule. Since $N2$ cannot be applied improperly, we focus on $N1$. We will show that applying an $N1$ rule improperly cannot increase the value. The desired result will then follow by induction.

We will argue that for any $z \in \mathbb{N}$ and any vector $\mathbf{c} \in \mathbb{N}^{k+1}$ such that $|\mathbf{c}|_\infty \leq z$, it holds that for all $2 \leq i \leq k+1$, applying any rule from $\{N1_j \mid 0 < j < i\}$ improperly to $\Phi(\mathbf{c}; z)$, yields a value smaller than $\Phi(\mathbf{c}; z)$. For the base case we consider $i = 2$. We need to show the property holds for $N1_1$. Assume that $(\mathbf{c}, z) \rightsquigarrow_{N1_1} (c_k, \dots, c_1 - 1, z + 1, z)$ and that $c_0 > 0$. By applying c_0 times the rule $N2$ to $\Phi(\mathbf{c}; z)$ we obtain $\Phi(c_k, \dots, c_1, 0; z + c_0)$. We can now properly apply $N1_1$ to the latter and obtain $\Phi(c_k, \dots, c_1 - 1, z + c_0; z + c_0)$. It follows from Lemma 8 that the claim holds for $i = 2$. To conclude, we show that if it holds for i then it must hold for $i + 1$. We only need to show the property holds for $N1_i$. Assume that $(\mathbf{c}, z) \rightsquigarrow_{N1_i} (c_k, \dots, c_i - 1, z + 1, c_{i-2}, \dots, z)$ and that $c_\ell > 0$ for some $0 \leq \ell < i$. By applying the sequence of rules $N1_\ell N1_{\ell-1} \dots N1_1 N2 N1_i N1_{i-1} \dots N1_1$ to $\Phi(\mathbf{c}; z)$ we obtain $\Phi(c_k, \dots, c_i - 1, z + 1, z + 1, \dots, z + 2; z + 1)$. It follows from induction hypothesis and the fact that $|\mathbf{b}|_\infty \leq z$ that $\Phi(\mathbf{c}; z)$ is larger than the latter, which is, in turn, larger than $\Phi(c_k, \dots, c_i - 1, z + 1, c_{i-2}, \dots; z)$ according to Lemma 8. Thus, the claim holds. \square

In the next section we will detail a new gadget which can be used to pump an energy level of m up to $F_m(m) = F_\omega(m)$. The gadget simulates the rewrite rules to compute F_ω vectorially.

4.3 F_ω -pumping gadget

For convenience, we will focus on the blind gadget as a blind energy game itself. We will later comment on how it fits together with the Minsky machine game constructed in Section 4.1.

Let us consider a fixed $m \in \mathbb{N}$. The blind energy game I_m we build has exactly $m + 5$ states, namely: an initial state q_0 and states $\{\top, f, \chi\} \cup \{\alpha_i \mid 0 \leq i \leq m\}$. The game starts with a non-deterministic choice of state from the set $\{\chi\} \cup \{\alpha_i \mid 0 \leq i \leq m\}$. The weight of the transition going to state α_m is 1; to state χ , m ; to all other states, 0. The alphabet consists of as many rewrite rules (defined in the previous section) as required to compute F_m . More formally, the alphabet is $\Sigma = \{N0, N2\} \cup \{N1_j \mid 1 \leq j \leq m\}$. For clarity, the game has been split into Figures 6–8, each Figure showing transitions for different letters. Intuitively, the game I_m allows Eve to simulate the vectorial computation of $F_\omega(m)$. Once she plays the letter N_0 the play moves to either \top or f , and—as we will argue later—if she has correctly simulated the computation of the function and the play has reached f , its energy level will be $F_\omega(m)$.

Let us write Σ_N for the restricted alphabet $\{N2\} \cup \{N1_j \mid 1 \leq j \leq m\} \subseteq \Sigma$ and Q_N for the set of states $\{\chi\} \cup \{\alpha_i \mid 0 \leq i \leq m\}$. We now prove the property this game (or gadget) enforces. The idea is that Eve playing a sequence of rewrite rules Σ_N has the effect that all plays consistent with her strategy and which end in α_i have energy level equal to the value of α_i after applying the rules to $\Phi(\mathbf{a}; m)$.

Lemma 12. *Consider any play prefix $\pi = q_0 \sigma_0 \dots \sigma_{n-1} q_n$ in I_m such that $\sigma_i \in \Sigma_N$ for all $0 < i < n$, and $\alpha_i \in Q_N$ and $\text{EL}(\pi[0..i]) \geq 0$ for all $0 < i \leq n$. If $q_n = \alpha_j$ then $\text{EL}(\pi) = a_j^n$, and if $q_n = \chi$ then*

$$\text{EL}(\pi) = m^n, \text{ where } \mathbf{a}^0 = (1, \overbrace{0, \dots, 0}^{m \text{ times}}), m^0 = m, \text{ and } (a_m^0, \dots, a_0^0, m^0) \rightsquigarrow_{\sigma_0} \dots \rightsquigarrow_{\sigma_{n-1}} (a_m^n, \dots, a_0^n, m^n).$$




Figure 7: Pumping gadget with only transitions for $\sigma = N1_j$ shown for states $\{\chi\} \cup \{\alpha_i \mid 0 \leq i \leq m\}$.




Figure 8: Pumping gadget with only transitions for $\sigma = N2$ shown for states $\{\chi\} \cup \{\alpha_i \mid 0 \leq i \leq m\}$.

Proof. We proceed by induction on n . Note that if $n = 1$, i.e. the play prefix has only two states, the energy level will be equal to 1 if the second state in the play is α_m , m if it is χ , and 0 otherwise. Hence the claim holds for prefixes of some length n . Let us argue that this also holds for prefixes of length $n + 1$. Consider an arbitrary play prefix $\pi = q_0\sigma_0 \dots \sigma_n q_{n+1}$ for which all assumptions hold. By induction hypothesis we have that $\text{EL}(\pi[0..n]) = a_j^n$ if $q_n = \alpha_j$ and m^n otherwise. If $\sigma_n = N2$, following the transitions shown in Figure 8 we get that: if $q_n = \alpha_j$, and $0 < j \leq m$ then the claim holds since $q_{n+1} = \alpha_j$ and $\text{EL}(\pi) = a_j^{n+1} = a_j^n$; also, if $q_n = \alpha_0$ the claim holds since $q_{n+1} = \alpha_0$ and $\text{EL}(\pi) = a_0^{n+1} = a_0^n - 1$ as expected; finally, if $q_n = \chi$ then it holds since $q_{n+1} = \chi$ and $\text{EL}(\pi) = m^{n+1} = m^n + 1$. Otherwise, if $\sigma_n = N1_\ell$, following the transitions shown in Figure 7 we get that: if $q_n = \alpha_j$, and $j \notin \{\ell, \ell - 1\}$ then the claim holds since $q_{n+1} = \alpha_j$ and $\text{EL}(\pi) = a_j^{n+1} = a_j^n$; if $q_n = \chi$ then it holds since either $q_{n+1} = \chi$ and $\text{EL}(\pi) = m^{n+1} = m^n$ or $q_{n+1} = \alpha_{\ell-1}$ and $\text{EL}(\pi) = a_{\ell-1}^{n+1} = m^n + 1$; it cannot be the case that $q_n = \alpha_{\ell-1}$ since otherwise q_{n+1} must be \perp and that would violate our assumptions; finally, if $q_n = \alpha_\ell$ then $q_{n+1} = \alpha_\ell$ and $\text{EL}(\pi) = a_\ell^{n+1} = a_\ell^n - 1$ as required. The result thus follows by induction. \square

We are finally ready to prove our main result.

Theorem 2. *The fixed initial credit problem is ACK-hard, even for blind games.*

Proof. For any 2CM M we will consider two instances of the new pumping gadget I_m , with $m = |M|$, and one instance of the 2CM-simulating game G_M from Section 4.1. (The first copy of I_m will be used to compute the Ackermann function while the second one will be used to obtain a value greater than $m(F_\omega(m))^2$.) Additionally, we will add two new states, s_0 and s_1 , which have a 0-weighted self-loop on all letters from the alphabet of I_m , except for $N0$, and a bad sink state, \perp , which has self-loops with weight -1 on all letters from the alphabets of I_m and G_M . The good sink \top , in copies of I_m now have self-loops with weight 0 on all letters from the alphabet of G_M and I_m .

We describe how all five components are connected (see Figure 9). From χ in the first copy of I_m with $N0$ and weight 0 we non-deterministically go to s_0 and the initial state of the second I_m ; from s_0 with $N0$ and weight 0 we deterministically go to s_1 ; from s_1 with $N0$ and weight 0 we deterministically




Figure 9: Overview of the blind game used to show ACK-hardness.

go to the G_M gadget from Figure 4; from χ in the second copy of I_m with $N0$ and weight 0 we non-deterministically go to all copies of the gadgets from Figures 1, 2, 3, and 5. Finally, to make sure the transition relation is total, from both copies of I_m , s_0 , and s_1 we add transitions to \perp on all letters from the alphabet of G_M . Also, from G_M we add transitions to \perp on all letters from the alphabet of I_m .

We will now argue that Eve has an observation-based winning strategy for initial credit $c_0 = 0$ if and only if M has an $F_\omega(|M|)$ -bounded halting run.

If it halts, she wins If M has an $F_\omega(|M|)$ -bounded halting run, then Eve should play the sequence of proper rewrite rules required to compute $m' = F_\omega(|M|)$ vectorially from $\Phi(1, 0, \dots, 0; m)$. She will then play $N0$ and choose letters to compute $F_{m'+1}^{m'}(F_{m'-1}^{m'}(\dots F_0^{m'}(m' + m)))$ which we denote by m'' . Note that $m'' \geq |M|(F_\omega(|M|))^2$. Finally, she will simulate the halting run of M . From Lemma 12 we have that: after the first time she plays $N0$ the play is either in the first copy of \top (and will never have negative energy level) or it has energy level $F_\omega(|M|)$ and is now in the initial state of the second copy of I_m or s_0 . By playing $N0$ the play now moves to s_1 or the simulation of the computation of m'' . After playing a third $N0$, the play moves from the second copy of χ —with energy level m'' —and from s_1 —with m' —to the corresponding gadgets in G_M or it moves to \top , where Eve cannot lose. Hence, any play consistent with this strategy of Eve, and which does not enter G_M , cannot have negative energy level. If the play has entered G_M then the same arguments as presented to prove Proposition 2 should convince the reader that it cannot have negative energy level.

If it does not halt, she does not win If M has no $F_\omega(|M|)$ -bounded halting run then, from Lemmas 9 and 12 we have that Eve eventually play three times $N0$ to exit the copies of I_m and enter G_M —or end in a \top state—lest we can construct a play with negative energy level. Also, using Lemma 11, we conclude that she cannot exit the two copies of I_m and enter G_M with energy level greater than $F_\omega(|M|)$ for the gadget from Figure 4 or energy level greater than m'' for the other gadgets, respectively. It thus follows from the proof of Proposition 2 that she has no observation-based winning strategy. \square

References

- [AG11] Krzysztof R Apt and Erich Grädel. *Lectures in game theory for computer scientists*. Cambridge University Press, 2011.
- [BFL⁺08] Patricia Bouyer, Ulrich Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jirí Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *FORMATS*, volume 5215 of *LNCS*, pages 33–47. Springer, 2008.
- [CdAHS03] Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource interfaces. In Rajeev Alur and Insup Lee, editors, *EMSOFT*, volume 2855 of *LNCS*, pages 117–133. Springer, 2003.
- [CDHR10] Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Generalized mean-payoff and energy games. In *FSTTCS*, pages 505–516, 2010.

- [DDG⁺10] Aldric Degorre, Laurent Doyen, Raffaella Gentilini, Jean-François Raskin, and Szymon Toruńczyk. Energy and mean-payoff games with imperfect information. In *CSL*, pages 260–274, 2010.
- [FFSS11] Diego Figueira, Santiago Figueira, Sylvain Schmitz, and Philippe Schnoebelen. Ackermannian and primitive-recursive bounds with dickson’s lemma. In *LICS*, pages 269–278. IEEE Computer Society, 2011.
- [FW98] Matt Fairtlough and Stanley S Wainer. Hierarchies of provably recursive functions. *Handbook of proof theory*, 137:149–207, 1998.
- [HPR14] Paul Hunter, Guillermo A. Pérez, and Jean-François Raskin. Mean-payoff games with partial-observation - (extended abstract). In *RP*, pages 163–175, 2014.
- [JLS15] Marcin Jurdziński, Ranko Lazic, and Sylvain Schmitz. Fixed-dimensional energy games are in pseudo-polynomial time. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *ICALP*, volume 9135 of *LNCS*, pages 260–272. Springer, 2015.
- [Jur98] Marcin Jurdziński. Deciding the winner in parity games is in $UP \cap coUP$. *IPL*, 68(3):119–124, 1998.
- [LW70] Martin H Löb and Stanley S Wainer. Hierarchies of number-theoretic functions. i. *Archive for Mathematical Logic*, 13(1):39–51, 1970.
- [Sch16] Sylvain Schmitz. Complexity hierarchies beyond elementary. *TOCT*, 8(1):3, 2016.
- [SS12] Sylvain Schmitz and Philippe Schnoebelen. Algorithmic aspects of wqo theory. Lecture Notes, 2012.
- [Wai70] Stanley S Wainer. A classification of the ordinal recursive functions. *Archive for Mathematical Logic*, 13(3):136–153, 1970.
- [ZP96] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *TCS*, 158(1):343–359, 1996.